

Cognitive Indoor Positioning and Tracking using Multipath Channel Information

Erik Leitinger, *Member, IEEE*, Simon Haykin, *Life Fellow, IEEE*, and Klaus Witrisal, *Member, IEEE*

Abstract—This paper presents a *robust and accurate* positioning system that adapts its behavior to the surrounding environment like the visual brain, mimicking its capability of filtering out clutter and focusing attention on activity and relevant information. Especially in indoor environments, which are characterized by harsh multipath propagation, it is still elusive to achieve the needed level of accuracy *robustly* under the constraint of reasonable infrastructural needs. In such environments it is essential to separate relevant from irrelevant information and attain an appropriate uncertainty model for measurements that are used for positioning.

Index Terms—Cognitive dynamic systems, Cramér-Rao bounds, localization, channel models, ultra wideband communication

I. INTRODUCTION

A. Motivation and State of the Art

For radiobased positioning in indoor environments, which are characterized by harsh multipath propagation, it is still elusive to achieve the needed level of accuracy *robustly*¹ under the constraint of reasonable infrastructural needs. In such environments it is essential to separate relevant from irrelevant information and attain an appropriate uncertainty model for measurements that are used for positioning. A system that has the potential to overcome the impairments of indoor environments has to be able to adapt its behavior to the surrounding environment like the visual brain, mimicking its capability of filtering out clutter and focusing attention on activity and relevant information.

To approach this objective more closely the four basic principles for *human cognition*, namely the *perception-action-cycle (PAC)*, *memory*, *attention* and *intelligence* [1] are implemented into the positioning systems as schematically illustrated in Fig. 2. To encounter all these principles, the concepts of multipath-assisted indoor navigation and tracking (MINT) [2]–[4] are intertwined with the principles of cognitive dynamic systems (CDS) that were developed by Simon Haykin and co-workers [5]–[9]. Evidently, a perceptive system has to reason with measurements under *uncertainty* [10], i.e. it has to treat the gained information *probabilistically* [11], [12], but it also has to deliberately take actions on the environment and consequently influence measurements to reason in favor of

relevant information instead of irrelevant one. Hence, cognitive processing of measurement data for positioning seems to be a natural choice to overcome such severe impairments.

MINT exploits specular multipath components (MPCs) (by black lines) that can be associated to the local geometry as illustrated in Fig. 1 [2]–[4]. MPCs can be seen as signals from additional virtual sources, so-called virtual anchors (VA), that are mirror-images of a physical anchor w.r.t. the floor plan (blue square-crosses in Fig. 1) [2], [13], [14]. Ultra-wideband (UWB) signals are used because of their superior time resolution and to facilitate the separation of MPCs. Hence, additional *position-related information* is exploited that is contained in the radio signals. For a proper consideration of uncertainties in floor plan and stochastic nature of the radio signals a geometry-based stochastic channel model (GSCM) and geometry-based probabilistic environment model (GPEM) were introduced in [15] to extend MINT to a simultaneous localization and mapping (SLAM) approach. Such a system acquires and adapts online information about its surrounding environment and is able to continuously building up a consistent memory in a Bayesian sense.

The idea of combining MINT with a CDS is to gain control over the observed environment information to (i) provide as much position-related information to the Bayesian state estimator as possible for achieving the highest level of reliability/robustness in position estimation, (ii) to improve the separation between relevant and irrelevant information, and (iii) building up a consistent environment and action memory. By additionally actively planning next control actions on the environment using the Bayesian memory—in sense of waveform adaptation [5], [16]–[18] or mobile agent motor-control [19], [20]—the relevant information-return contained in the signals can be maximized. The information-flow coupling between the perceptor-actor system and the surrounding environment is given by the PAC that plays the key-role when it is coming to gather relevant environment information [1], [9].

The core feedback loop of the CDS, the perception-action-cycle resembles the idea of optimally choosing future measurements based on a physical model under reasoning with uncertainty. The principle has been explored by the physics community under the term Bayesian experimental design [21]. This decision-theoretic process gives a mathematical justification for selecting the appropriate optimality criterion under uncertainty that maximizes the utility function of the posterior probability density function, such that new model information of the acquired measurements can be predicted. Information theoretic measures as the conditional entropy [22], the mutual

E. Leitinger, and K. Witrisal are with Graz University of Technology, Graz, Austria, email: {erik.leitinger, paul.meissner, witrisal}@tugraz.at and Simon Haykin is with the Cognitive Systems Laboratory, McMaster University {haykin@mcmaster.ca}

This work was supported by the Austrian Research Promotion Agency (FFG) within the project REFlex (project number: 845630).

¹We define robustness as the percentage of cases in which a system can achieve its given potential accuracy.

information [22] or the determinate of the Fisher information matrix [23], [24] are suitable utility functions for this process. The active selection of measurement parameters has a lot in common with cognitive perception and control at the lowest layer. However, it lacks an explicit description of a layered memory structure that leads in combination with algorithmic attention to an “intelligent” behavior of the overall system.

In robotics, predictive optimal control is used to actively plan a robot’s path for achieving certain goals under constraints [25], [26] or for efficiently building better maps of environments in simultaneous localization and mapping (SLAM) [27]–[30]. In [31] decentralized information-theoretic control of autonomous systems has been analyzed with multiple sensors. Specifically, the predictive coordination and cooperation of systems has been advanced. In a similar manner, in [32], an information-theoretic sensor management approach has been presented for predicting configurations and sets of sensors. In [33], [34] efficient non-parametric reinforcement learning algorithm based on Gaussian processes has been developed and used for non-linear control. With probabilistic Gaussian processes, model uncertainties can be explicitly incorporated into the learning and planning process and one does not rely on specific state transition and measurement models, but rather on a set of functions that depend on the underlying uncertainty.

II. GEOMETRY-BASED PROBABILISTIC ENVIRONMENT MODEL (GPEM)

The basic work on MINT can be found in [2], [3], [35], [36]. In [2], a detailed mathematical description of how reflectors can be represented using virtual sources and how their positions are computed using optical ray-tracing with a priori knowledge of a floor plan.

A. Floor Plan with uncertainties – Probabilistic VA Positions

In the case, the floor plan is exactly known, the VA positions $\mathcal{A}_n^{(j)} = \{\mathbf{a}_{k,n}^{(j)}\}_{k=1}^{K^{(j)}}$ are computed by mirroring the j -th anchor w.r.t. walls, first-order VA are constructed. Higher-order VAs are computed by mirroring VAs w.r.t. walls [35] and [13], [14]. The VAs of all anchors are comprised in $\mathcal{A}_n = \{\mathcal{A}_n^{(j)}\}_j$. To be able to cope with uncertainties in the floor plan the deterministic geometric model of the VA positions $\mathbf{a}_k^{(j)}$ of the j -th anchor, is extended to a probabilistic one as shown in Fig. 1. The VA positions and the agent position $\mathbf{p}^{(m)}$ are represented by a joint PDF $p(\mathbf{p}^{(m)}, \mathbf{a}_1^{(j)}, \mathbf{a}_2^{(j)}, \dots, \mathbf{a}_{K^{(j)}}^{(j)})$. If the position of the j -th anchor is assumed to be known exactly, the joint PDF reduces to $p(\mathbf{p}^{(m)}, \mathbf{a}_2^{(j)}, \dots, \mathbf{a}_{K^{(j)}}^{(j)})$.

Fig. 1 illustrates the probabilistic geometric environment model. A signal exchanged between an anchor at position $\mathbf{a}_1^{(j)}$ and an agent at $\mathbf{p}^{(m)}$ contains specular reflections at the room walls, indicated by the black lines². These reflections can be modeled geometrically using the VA $\mathbf{a}_k^{(j)}$ with $k = 1, \dots, K^{(j)}$ that are mirror-images of the j -th anchor w.r.t. walls [2], [13], [14]. The number of expected VA per anchor j is defined as $K^{(j)}$.

²Since the radio channel is reciprocal, the assignment of transmitter and receiver roles to anchors and agents is arbitrary and this choice can be made according to the application scenario.




Fig. 1. Illustration of the VA for the j -th anchor and an agent with PDF $p(\mathbf{a}_k^{(j)})$ and $p(\mathbf{p}^{(m)})$, respectively. The VA at position $\mathbf{a}_{\text{false}}$ represents a false detected VA.

The joint PDF of the agent and the VA positions is represented by a multivariate Gaussian RV, where the figure shows the marginal distributions of the agent $p(\mathbf{p}^{(m)})$ (dashed black ellipses) and the VA positions $p(\mathbf{a}_k^{(j)})$ (red ellipses). The marginal distribution $p(\mathbf{a}_{\text{false}})$ (dashed red ellipse) defines a wrongly detected VA at position $\mathbf{a}_{\text{false}}$. The anchor position $\mathbf{a}_1^{(j)}$ is assumed to be known perfectly. Uncertainty in the floor plan does not just mean that the VA positions are uncertain and thus described by RV, but also that floor plan information is incorrect/inconsistent or entirely missing. This means that positioning and tracking algorithms based on VA, have to consider this lack of knowledge.

III. GEOMETRY-BASED STOCHASTIC CHANNEL MODEL (GSCM)

In this section we review the basic concept of MINT starting with the signal model, then discussing the estimation of the MPC parameters, and finally introducing position related information that is of main importance for a proper weighting of the MPC-VA relations in the Bayesian tracking filter. All not-geometrically-modeled propagation effects in the signals constitute interference to the useful position-related information, so-called diffuse multipath (DM) [37]. An online estimation of the influence of the DM on the range uncertainties to the VAs allows for an efficient selection of the VAs that can be reliably used for the agent tracking [15].

In this section we review the basic concept of MINT starting with the signal model, then discussing the estimation of the MPC parameters, and finally introducing position related information that is of main importance for a proper weighting of the MPC-VA relations in the Bayesian tracking filter. All not-geometrically-modeled propagation effects in the signals constitute interference to the useful position-related information, so-called diffuse multipath (DM) [37]. An online estimation of the influence of the DM on the range uncertainties to the VAs allows for an efficient selection of the VAs that can be reliably used for the agent tracking [15].

A. Signal Model

During time step n , a baseband UWB signal $s(t)$ with effective pulse duration T_p is transmitted from the j -th anchor

located at position $\mathbf{a}_1^{(j)} \in \mathbb{R}^2$, with $j \in \{1, \dots, J\}$, to an agent at position $\mathbf{p}_n \in \mathbb{R}^2$. The corresponding received signal is modeled as [3]

$$r_n^{(j)}(t) = \sum_{k=1}^{K_n^{(j)}} \alpha_{k,n}^{(j)} s(t - \tau_{k,n}^{(j)}) + (s * \nu_n^{(j)})(t) + w(t). \quad (1)$$

The first term describes the sum of $K_n^{(j)}$ deterministic MPCs with complex amplitudes $\alpha_{k,n}^{(j)}$ and delays $\tau_{k,n}^{(j)}$, where $k \in \{1, \dots, K_n^{(j)}\}$. The delays correspond to the distances to VAs of the j -th anchor at positions $\mathbf{a}_k^{(j)} \in \mathbb{R}^2$, $k \in \{1, \dots, K_n^{(j)}\}$, i.e., $\tau_{k,n}^{(j)} = \frac{1}{c} \|\mathbf{p}_n - \mathbf{a}_k^{(j)}\|$, where c is the speed of light. The delay $\tau_{1,n}^{(j)}$ corresponds to the direct line-of-sight path between the j -th anchor and the mobile agent. The energy of signal $s(t)$ is assumed to be normalized to one. The second term denotes the convolution of the transmitted signal $s(t)$ with the diffuse multipath function $\nu_n^{(j)}(t)$, which is modeled as a non-stationary zero-mean Gaussian random process. For the diffuse multipath, we assume uncorrelated scattering along the delay axis τ , hence the auto-correlation function of $\nu_n^{(j)}(t)$ is given by $\mathbb{E}_{\nu} \{ \nu_n^{(j)}(\tau) \nu_n^{(j)*}(u) \} = S_{\nu,n}^{(j)}(\tau) \delta(\tau - u)$, where $S_{\nu,n}^{(j)}(\tau)$ is the power delay profile of the diffuse multipath at the agent position \mathbf{p}_n . The diffuse multipath process is assumed to be quasi-stationary in the spatial domain, which means that $S_{\nu,n}^{(j)}(\tau)$ does not change in the vicinity of \mathbf{p}_n [38]. Note that the diffuse multipath component infers with the useful position-related information. The last term in (1), $w(t)$, is additive white Gaussian noise with double-sided power spectral density $N_0/2$.

B. MPC Parameter Estimation

The MPC delays at agent position \mathbf{p}_n are estimated recursively by a least-squares approximation of the received signal [35]. The estimate of the m -th MPC delay is calculated as

$$\hat{\tau}_{m,n}^{(j)} = \arg \min_{\tau} \int_0^T |r_n^{(j)}(t) - \hat{r}_{n,m-1}^{(j)}(t) - \hat{\alpha}(\tau) s(t - \tau)|^2 dt, \quad (2)$$

where T is the measurement duration, $\hat{r}_{n,m-1}^{(j)}(t)$ is the template signal containing all MPCs up to the $(m-1)$ -th, and

$$\hat{\alpha}(\tau) = \int_0^T [r_n^{(j)}(t) - \hat{r}_{n,m-1}^{(j)}(t)] s^*(t - \tau) dt. \quad (3)$$

Furthermore, $\hat{\alpha}_{m,n}^{(j)}$ is an estimate of the m -th amplitude that is obtained as $\hat{\alpha}_{m,n}^{(j)} = \hat{\alpha}(\hat{\tau}_{m,n}^{(j)})$. The template signal is formed as $\hat{r}_{n,m-1}^{(j)}(t) = \sum_{m'=1}^{m-1} \hat{\alpha}_{m',n}^{(j)} s(t - \hat{\tau}_{m',n}^{(j)})$. Alternating between (2) and (3), the algorithm recursively estimates the MPC parameters $\tau_{m,n}^{(j)}$ and $\alpha_{m,n}^{(j)}$ until a predefined number M of MPCs is reached. The template signal is initialized with $\hat{r}_{n,0}^{(j)}(t) = 0$.

The estimated delays are scaled by the speed of light c , i.e., $y_{m,n}^{(j)} = c \hat{\tau}_{m,n}^{(j)}$, with $m \in \{1, \dots, M\}$, and used as noisy distance measurements. Furthermore, in MINT system, the amplitude estimates $\hat{\alpha}_{m,n}^{(j)}$ (after being associated with the k -th VA) are fed into a higher-level, non-Bayesian algorithm that determines the signal-to-interference-plus-noise power

ratio between the useful deterministic MPC and the diffuse-multipath plus noise. This power ratio is a measure of the variance $\sigma_{k,n}^{(j)2}$ of the measurement $y_{k,n}^{(j)}$ that corresponds to the k -th VA (see [3], [35] for details). The finite set of measured delays is written as $\mathcal{Z}_n = \bigcup_j \mathcal{Z}_n^{(j)} = \bigcup_j \{ \hat{d}_{k,n}^{(j)} \}_{k=1}^{K_n^{(j)}}$.

C. Position and Range Uncertainty

As a performance measure and lower bound we use the Cramer-Rao-Lower Bound (CRLB) of the position error defined by the inequality $\mathbb{E} \{ \|\mathbf{p} - \hat{\mathbf{p}}\|_2^2 \} \geq \text{tr} \{ \mathbf{J}_{\mathbf{p}}^{-1} \}$, where $\mathbf{J}_{\mathbf{p}}$ is the equivalent Fisher information matrix (EFIM) [3], [39], [40] for the position vector and $\text{tr} \{ \cdot \}$ is the trace operator. Assuming no path overlap between MPCs, the EFIM $\mathbf{J}_{\mathbf{p}}$ is formulated for a set of anchors in a canonical form by [3]

$$\mathcal{I}_{\mathbf{p},n} = \frac{8\pi^2 \beta^2}{c^2} \sum_{j=1}^J \sum_{k=1}^{K_n^{(j)}} \text{SINR}_{k,n}^{(j)} \mathcal{I}_{\mathbf{r}}(\phi_{k,n}^{(j)}), \quad (4)$$

where β denotes the effective (root mean square) bandwidth of $s(t)$ and $\mathcal{I}_{\mathbf{r}}(\phi_{k,n}^{(j)})$ the ranging direction matrix, which has an eigenvector in direction $\phi_{k,n}^{(j)}$ from the agent to the k -th VA. The SINRs are described by the relation between the energy of the deterministic MPCs to the power of DM plus noise

$$\text{SINR}_{k,n}^{(j)} = \frac{|\alpha_{k,n}^{(j)}|^2}{N_0 + T_{\mathbf{p}} S_{\nu,n}^{(j)}(\tau_{k,n}^{(j)})} \quad (5)$$

The according MPC range uncertainties $\sigma_{d,k,n}^{(j)2} = \text{var} \{ \hat{d}_{k,n}^{(j)} \}$ to already associated VAs is given as

$$\sigma_{d,k,n}^{(j)2} \geq \left(\frac{8\pi^2 \beta^2}{c^2} \text{SINR}_{k,n}^{(j)} \right)^{-1}. \quad (6)$$

The SINRs are estimated with a method of moments estimator [35] using the associated complex amplitudes $\{ \hat{\alpha}_{k,i}^{(j)} \}_{i=n-n_0}^n$ over a window of past agent positions where the MPC was associated. The online estimation is started once an initial window size of measurements is available a VA. Until then, a default value $\sigma_{d,\text{init}}^2$ is assigned.

D. Probabilistic Data Association (PDA)

The set of expected MPC delays $\mathcal{D}_n^{(j)}$ at time step n is computed as the distances of each VA in $\mathcal{A}_n^{(j)}$ to the predicted position of the mobile agent is $\mathbf{x}_n = [\mathbf{p}_n^T, \mathbf{v}_n^T]^T$, where \mathbf{v}_n is the velocity. The state evolves according to the state transition probability density function (PDF) $p(\mathbf{x}_n | \mathbf{x}_{n-1})$ over time instances n . The MPC distances described in Section III-B are subject to a data association uncertainty, i.e., it is not known which measurement in $\mathbf{y}_n^{(j)}$ originated from which VA k of the j -th anchor, and it is also possible that a measurement $y_{m,n}^{(j)}$ did not originate from any VA (false alarm, clutter) or that a VA did not give rise to any measurement (miss detection). The probability that a VA is detected is denoted by P_d . Possible associations at time instance n are described by the $K_n^{(j)}$ -dimensional random vector $\mathbf{b}_n^{(j)} = [b_{1,n}^{(j)} \dots b_{K_n^{(j)},n}^{(j)}]^T$, whose

k -th entry is defined as

$$b_{k,n}^{(j)} = \begin{cases} m \in \{1, \dots, M\}, & \mathbf{a}_k^{(j)} \text{ generates measurement } y_{m,n}^{(j)} \\ 0, & \mathbf{a}_k^{(j)} \text{ did not give rise to any measurement.} \end{cases}$$

We also define $\mathbf{b}_n = [\mathbf{b}_n^1 \dots \mathbf{b}_n^J]^T$. False alarms are modeled by a uniform distribution with mean arrival rate μ , and the distribution of each false alarm measurement is described by the PDF $f_{\text{FA}}(y_{m,n}^{(j)})$ [41], [42].

The statistical dependence of the distance measurement vectors $\mathbf{y}_n = [\mathbf{y}_n^1, \dots, \mathbf{y}_n^J]^T$ on the agent state vector \mathbf{x}_n and the association vector \mathbf{b}_n is described by the *global likelihood function* $f(\mathbf{y}_n | \mathbf{x}_n, \mathbf{b}_n)$. Under commonly-used assumptions about the statistics of the measurements [2], [41], [42], the global likelihood function at time instances n factors as

$$f(\mathbf{y}_n | \mathbf{x}_n, \mathbf{b}_n) = \prod_{j=1}^J \left(\prod_{m=1}^M f_{\text{FA}}(y_{m,n}^{(j)}) \right) \times \prod_{k \in \mathcal{Q}(\mathbf{x}_n, \mathbf{b}_n^{(j)})} \frac{f(y_{b_{k,n}^{(j)}}^{(j)} | \mathbf{x}_n; \mathbf{a}_k^{(j)})}{f_{\text{FA}}(y_{b_{k,n}^{(j)}}^{(j)})},$$

where $\mathcal{Q}(\mathbf{x}_n, \mathbf{b}_n^{(j)}) \triangleq \{k \in \{1, \dots, K_n^{(j)}\} : b_{k,n}^{(j)} \neq 0\}$. Here, the local likelihood function $f(y_{b_{k,n}^{(j)}}^{(j)} | \mathbf{x}_n; \mathbf{a}_k^{(j)}, \sigma_{d,k,n}^{(j)})$ of $y_{m,n}^{(j)}$ is related to a noisy measurement of the distance to VA $\mathbf{a}_k^{(j)}$ at agent position \mathbf{p}_n modeled as

$$y_{k,n}^{(j)} = \|\mathbf{p}_n - \mathbf{a}_k^{(j)}\| + v_{k,n}^{(j)},$$

where $v_{k,n}$ is zero-mean Gaussian noise with variance $\sigma_{d,k,n}^{(j)2}$ as described in (6). Based on the factorized likelihood model, a probabilistic data association algorithm using belief propagation as described in [43], [44] is used to compute the associations between the expected delay to the VAs and the estimated MPCs. The most probable MPC-to-anchor association are obtained by means of an approximation of the maximum a posterior (MAP) detector [45]

$$\hat{c}_{k,n}^{(j)\text{MAP}} \triangleq \arg \max_{c_{k,n}^{(j)} \in \{1, \dots, M\}} p(c_{k,n}^{(j)} | \mathbf{y}). \quad (7)$$

After the PDA was applied for all anchors, the following union sets are defined:

- The set of associated discovered (and optionally a-priori known) VAs $\mathcal{A}_{n,\text{ass}} = \bigcup_j \mathcal{A}_{n,\text{ass}}^{(j)}$.
- The according set of associated measurements $\mathcal{Z}_{n,\text{ass}} = \bigcup_j \mathcal{Z}_{n,\text{ass}}^{(j)}$.
- The set of remaining measurements $\mathcal{Z}_{n,\text{ass}} = \bigcup_j \mathcal{Z}_{n,\text{ass}}^{(j)}$, which are not associated to VAs of $\mathcal{A}_{n,\text{ass}}$.

IV. COGNITIVE POSITIONING SYSTEM

The basic building blocks of a CDS are depicted in Fig. 2 and described as follows:

- *Cognitive perceptor (CP)*: Has the purpose to focus the attention on extracting and separating relevant information and irrelevant information from the received signal.

Irrelevant information is coupled with uncertainty and is therefore used as probabilistic weighting for relevant information to increase the robustness of a Bayesian state estimator. Further, a higher-level environment model is learned and memorized.

- *Probabilistic reasoning machine (PRM) and information feedback*: Mediates the information flow between the cognitive perceptor and the controller. The feedback information is described by an *information-theoretic measure* of the probabilistic state of the CP.
- *Cognitive controller (CC)*: Has the purpose to act on the environment based on the feedback information of the CP and to memorize past actions and to enable predictive planning.
- *Perception-action cycle (PAC)*: Incorporates the sensed environment into the closed loop with the CP and the CC. This means that every action on the environment directly influences the CP in the next sensing period.
- *Reciprocal Coupling*: As shown in Figure 2, the left and the right parts of the CDS are reciprocally coupled via the PRM so that a persistent information flow between both sides is guaranteed. This means that the feedback information of the CP is driving the cognitive actions, but also the CC is interacting with the CP via the PRM to guarantee consistency in the chosen cognitive actions by using the perceptive memory.
- *Hierarchical Structure*: Enables the ability to interpret the environmental observables on different abstraction layers. The level of abstraction is increased as we go up from layer to layer, so that the characteristic probabilistic essence of the observables is partitioned over the CDS and the *robustness* of the state estimation can be controlled. The layers of the CP are interacting with one another to continuously adapt the memory of past environment experiences to the changes in the environment in a joint manner. In the same manner also the CC has to have this hierarchical structure to actively sense the environment for every aspect of abstraction in an optimal way.

A. Multipath-assisted Positioning as CDS

Figure 2 illustrates the block diagram of a cognitive localization and tracking system with a triple layered structure:

- *First Layer*: Defines (i) the direct Bayesian state estimation $p(\mathbf{x}_n | \mathcal{Z}_n, \mathbf{b}_n^{(j)})$ at the CP, where $\mathbf{x}_n = [\mathbf{p}_n^T, \mathbf{v}_n^T]^T$ holds the agent position and its velocity, and (ii) the waveform adaptation at the CC based on the feedback information of the Bayesian state space filter.
- *Second Layer*: Represents (i) the memory for the GPEM described by the VA with marginal PDF $p(\mathcal{A}_n^{(j)} | \mathcal{Z}_n^{(j)}, \mathbf{b}_n^{(j)})$ and the memory for the GSCM described by the $\text{SINR}_{k,n}^{(j)}$ of the MPC at the CP and (ii) the memory of VA specific signal parameters c_n , e.g. the pulse duration T_p and/or carrier frequency f_c , at the CC.
- *Third Layer*: It represents the highest layer and is different from the two layers below in the sense that it

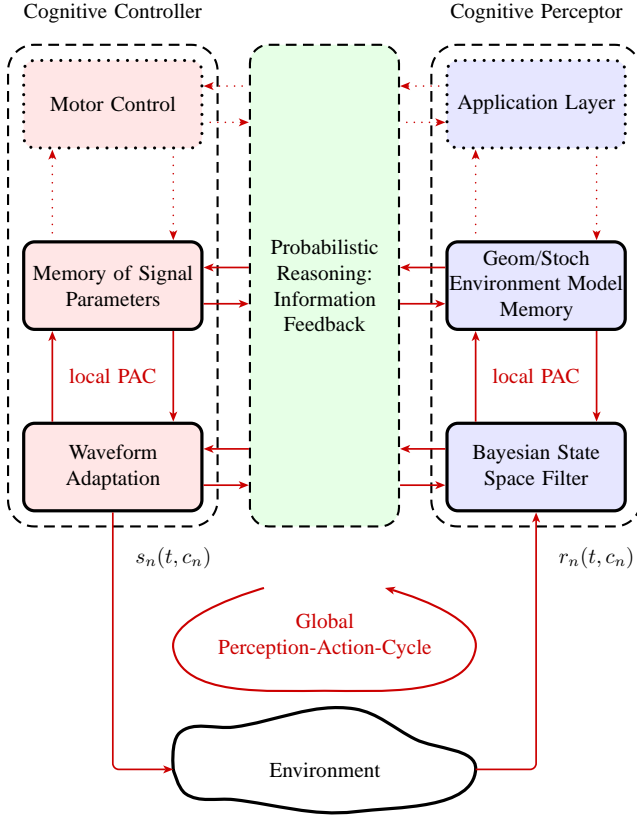


Fig. 2. Block diagram of Cognitive Dynamic System (CDS) for indoor positioning and tracking.

defines the application driven by the cognitive localization/tracking system. The CP memory of applications holds abstract parameters or structures of the specified application and the CC enables the motor control for realizing higher goal planning [46].

The first and second layers describe the signal and information processing of the model parameters of the surrounding physical environment and the radio channel. On the other hand, the third layer holds higher goal parameters, i.e. motor-control input to fulfill navigation goals, that are based on the physical-related parameters [31], [46], [47].

The control parameter vector \mathbf{c}_{n+1} of the next sensing cycle is chosen in order to gain the most “valuable” position-related information from the new set of measurements $\tilde{\mathbf{Z}}_{n+l}$ using the predicted posterior $p(\mathbf{x}_{n+l}, \mathcal{A}_{n+l} | \tilde{\mathbf{Z}}_{n+l}, \mathbf{b}_{n+l}, \mathbf{c}_{n+l})$, with $l = 0, \dots, l_{\text{future}}$ as future horizon. This goal can be reached by minimizing an expected cost-to-go function, yielding

$$\mathbf{c}_{n+1} = \arg \min_{\mathbf{c}_n} \mathcal{C}(p(\mathbf{x}_{n+l}, \mathcal{A}_{n+l} | \tilde{\mathbf{Z}}_{n+l}, \mathbf{b}_{n+l}, \mathbf{c}_n)), \quad (8)$$

where $\mathcal{C}(\cdot)$ is the expected cost-to-go function for optimal control [21], [31] of the environmental information contained in the predicted received signals $\tilde{\mathbf{Z}}_{n+l}$ that depends on the chosen signal model. The expected cost-to-go function is based on an information-theoretic measure that should depend on the environment parameters, like the VA specific $\text{SINR}_{k,n}^{(j)}$,

and serves as feedback information in the CDS. In general, estimation and control problems have to deal with probabilistic states and observations. As a consequence, also the control has to be probabilistic, i.e. the cost function or utility must handle uncertainties. Based on covering the uncertainty of the state with a PDF, a measure of informativeness of measurements has to be defined on the posterior state distribution. Two commonly used information measures of an RV are the entropy [48] and the Fisher information [24].

B. Information Measures

1) *Entropy*: For a continuous-valued vector RV $\mathbf{p} \in \mathbb{R}^L$ (in the follow-up sections \mathbf{p} represents the agent position), the conditional entropy is given as [22]

$$h(\mathbf{p}) \doteq -\mathbb{E}_{\mathbf{p}} \{\ln p(\mathbf{p})\} = - \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} p(\mathbf{p}) \ln p(\mathbf{p}) d\mathbf{p}, \quad (9)$$

The entropy is directly related to the uncertainty of the according RV. For a multivariate Gaussian RV $\mathcal{N}(\mathbf{m}_{\mathbf{p}}, \mathbf{C}_{\mathbf{p}})$ this means that the entropy is directly related to the covariance matrix $\mathbf{C}_{\mathbf{p}}$, yielding

$$h(\mathbf{p}) = \frac{1}{2} \ln ((2\pi e)^L \mathbf{C}_{\mathbf{p}})), \quad (10)$$

where $\det(\cdot)$ defines the determinant of a matrix. The connection to information one could see in the determinate of the covariance matrix $\mathbf{C}_{\mathbf{p}}$, which is a measure of the “volume” of the RV \mathbf{p} . The more compact the volume is, the smaller is the entropy $h(\mathbf{p})$ and consequently the more informative is the distribution $p(\mathbf{p})$.

2) *Fisher Information*: The Fisher information matrix (FIM) of a RV \mathbf{r} , dependent on the deterministic parameter \mathbf{p} , can also be used as a measure of information. Using the likelihood function $\ln f(\mathbf{r}; \mathbf{p})$, it is defined as

$$\mathcal{I}_{\mathbf{p}} = \mathbb{E}_{\mathbf{r}; \mathbf{p}} \left\{ \left[\frac{\partial}{\partial \mathbf{p}} \ln f(\mathbf{r}; \mathbf{p}) \right] \left[\frac{\partial}{\partial \mathbf{p}} \ln f(\mathbf{r}; \mathbf{p}) \right]^T \right\}. \quad (11)$$

The inverse of the FIM is a lower bound on the covariance $\mathbf{C}_{\hat{\mathbf{p}}} \succeq \mathcal{I}_{\mathbf{p}}^{-1}$ of the deterministic parameter \mathbf{p} of an estimator $\hat{\mathbf{p}}$ [24]. Looking at the entropy of the estimator’s distribution $\mathcal{N}(\hat{\mathbf{p}}, \mathbf{C}_{\hat{\mathbf{p}}})$, the explicit relationship between the FIM $\mathcal{I}_{\mathbf{p}}$ of \mathbf{r} (dependent on \mathbf{p}) and the entropy $h(\hat{\mathbf{p}})$ is given as

$$\begin{aligned} h(\hat{\mathbf{p}}) &= \frac{1}{2} \log ((2\pi e)^L \det (\mathbf{C}_{\hat{\mathbf{p}}})) \\ &\geq -\frac{1}{2} \log ((2\pi e)^L \det (\mathcal{I}_{\mathbf{p}})). \end{aligned} \quad (12)$$

As the relationship in (12) shows, one can connect the FIM of a parameter vector with the entropy, resulting in a scalar measure of information that is valuable for choosing optimal waveform parameters, as it is needed for a cognitive positioning system. As it is shown in Section III-C, the FIM $\mathcal{I}_{\mathbf{p}}$ on the position of the agent \mathbf{p} contains the environment and signal parameters, e.g. VA positions and the according SINRs. With this, a direct relationship between the environment, the feedback information and the control of the sensing is given, closing the PAC (Figure 2). In the same manner, the system

can also be expanded to information-based control of the agent state to increase the informativeness in the measurements [19], [20], [47].

V. COGNITIVE MINT

A. Bayesian State Space Estimator

In the most generic form, the prediction equation for the VAs \mathcal{A}_n and the agent state $\mathbf{x}_n = [\mathbf{p}_n, \mathbf{x}_n]^T$, can be written as, using the Markovian assumption,

$$p(\mathbf{x}_n, \mathcal{A}_n | \mathcal{Z}_{1:n-1}) = \int p(\mathbf{x}_{n-1}, \mathcal{A}_{n-1} | \mathcal{Z}_{1:n-1}) p(\mathbf{x}_n | \mathbf{x}_{n-1}) \times p(\mathcal{A}_n | \mathcal{A}_{n-1}) d\{\mathbf{x}_{n-1}, \mathcal{A}_{n-1}\}, \quad (13)$$

where $p(\mathbf{x}_n | \mathbf{x}_{n-1})$ and $p(\mathcal{A}_n | \mathcal{A}_{n-1})$ are the state transition probability distribution functions of the agent and the VAs, respectively. The latter can be represented by an identity function. The update equation is then

$$p(\mathbf{x}_n, \mathcal{A}_n | \mathcal{Z}_{1:n}) = \frac{p(\mathcal{Z}_n | \mathbf{x}_n, \mathcal{A}_n) p(\mathbf{x}_n, \mathcal{A}_n | \mathcal{Z}_{1:n-1})}{p(\mathcal{Z}_{1:n} | \mathcal{Z}_{1:n-1})}, \quad (14)$$

where $p(\mathcal{Z}_n | \mathbf{x}_n, \mathcal{A}_n)$ is the likelihood function of the current measurements. Assuming that the agent moves along a trajectories according to a linear Gaussian constant-velocity motion model can be written as,

$$\begin{aligned} \mathbf{x}_n &= \mathbf{F} \mathbf{x}_{n-1} + \mathbf{G} \mathbf{n}_{a,n} \\ &= \begin{bmatrix} 1 & 0 & \Delta T & 0 \\ 0 & 1 & 0 & \Delta T \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{x}_n + \begin{bmatrix} \frac{\Delta T^2}{2} & 0 \\ 0 & \frac{\Delta T^2}{2} \\ \Delta T & 0 \\ 0 & \Delta T \end{bmatrix} \mathbf{n}_{a,n}, \end{aligned} \quad (15)$$

where ΔT is the discrete time update rate. The driving acceleration noise term $\mathbf{n}_{a,n}$ is zero-mean, circular symmetric with variance σ_a^2 , and models motion changes which deviates from the constant-velocity assumption. The transformed noise covariance matrix is given as $\mathbf{R}_a = \sigma_a^2 \mathbf{G} \mathbf{G}^T$. The entire state space of \mathbf{x}_n and the associated VAs $\mathcal{A}_{n,\text{ass}}$ described in (13) is formulated as [4], [15]

$$\tilde{\mathbf{x}}_n = \begin{bmatrix} \mathbf{F} & \mathbf{0}_{4 \times 2K_n} \\ \mathbf{0}_{2K_n \times 4} & \mathbf{I}_{2K_n \times 2K_n} \end{bmatrix} \tilde{\mathbf{x}}_{n-1} + \begin{bmatrix} \mathbf{G} \\ \mathbf{0}_{2K_n \times 2} \end{bmatrix} \mathbf{n}_{a,n}, \quad (16)$$

where $\tilde{\mathbf{x}}_n = [\mathbf{x}_n^T, \mathbf{a}_{2,n}^T, \dots, \mathbf{a}_{K_n,n}^T]^T$ represents the stacked state vector, where $\{\mathbf{a}_{k,n}^{(j)}\} \in \mathcal{A}_{n,\text{ass}}$. The covariance matrix of the state vector consists of a the agent covariance matrix $\mathbf{C}_{\mathbf{x}_n}$, the cross-covariances between the agent state \mathbf{x}_n and the VAs at positions $\mathbf{a}_{k,n}$ $\mathbf{C}_{\mathbf{x}_n, \mathbf{a}_{k,n}}$, the cross-covariances between VAs $\mathbf{C}_{\mathbf{a}_{k,n}, \mathbf{a}_{k',n}}$ with $k \neq k'$, and the covariances of the VAs $\mathbf{C}_{\mathbf{a}_{k,n}}$. The measurement model is defined as

$$\tilde{\mathbf{z}}_n = \tilde{\mathbf{h}}_n(\tilde{\mathbf{x}}_n) + \tilde{\mathbf{n}}_{z,n} = \tilde{\mathbf{H}}_n \tilde{\mathbf{x}}_n + \tilde{\mathbf{n}}_{z,n}, \quad (17)$$

where $\tilde{\mathbf{z}}_n = [\mathbf{z}_n^T, \mathbf{z}_{2,\mathcal{P}_{2,n}}^T, \dots, \mathbf{z}_{K_n,\mathcal{P}_{K_n,n}}^T]^T$ is the stacked measurement vector. The vector $\mathbf{z}_{k,\mathcal{P}_{k,n}}$ represents a set of measurements over time (past measurements) that have been associated with the k -th VA, where $\mathcal{P}_{k,n}$ defines the according set of time indices at which the DA was possible. The stack vector $\tilde{\mathbf{n}}_{z,n}$ contains the according measurement noise. The measurement model $\tilde{\mathbf{h}}_n$ contains all distance equations from

the VA positions to the agent position $d(\mathbf{a}_{n,k}^{(j)}, \mathbf{p}_n) \forall \mathbf{a}_{n,k}^{(j)} \in \mathcal{A}_{n,\text{ass}}$ and from the actual and past agent positions to the VA positions $d(\mathbf{p}_n, \mathbf{a}_{n,k}^{(j)}) \forall \mathbf{a}_{n,k}^{(j)} \in \mathcal{A}_{n,\text{ass}}$ to update the agent and the VAs, respectively. As Bayesian state estimator a UKF is used [4]. The measurement covariance matrix is written as

$$\mathbf{R}_n = \text{diag} \left\{ \text{var} \left\{ \hat{d}_{k,n}^{(j)} \right\} \right\} \quad \forall k, j : \mathbf{p}_{k,n}^{(j)} \in \mathcal{A}_{n,\text{ass}}, \quad (18)$$

where the range variances are defined by (6).

B. Feedback Information

Fig. 3 illustrates the schematics of this cognitive tracking/localization algorithm. The system is able to adapt online its behavior to the environment, i.e. perceptual attention is given, through the following principles:

- At the CP side, the GSCM and GPDM memories are updated using the received signal $r_n(t, \mathbf{c}_n)$ with waveform parameters chosen by the CC.
- In the actual sensing cycle the attention is put through the CC using the control parameters \mathbf{c}_n on the potential set of VA and their parameters memorized in the GSCM and GPDM. These model parameters are seen at the CP side of Fig. 3.

Now the question is, “How to control the environment information flow through the received signal and put cognitive attention on the relevant features in the following sensing cycle?” The answer to this lies in the CC and the feedback and feed-forward information between the perceptor and the controller as illustrated in Fig. 3. The cognitive control algorithm introduced in this section is based on [7]. In there, the derivation of the so-called cognitive RL algorithm and the proof that the algorithm obtains Bellman optimality are presented.

As already stated in Section IV, the control parameters should be chosen in order to optimize the expected cost-to-go function $\mathcal{C}(\cdot)$ of the predicted posterior PDF as defined in (8). In general, the expected cost-to-go function for a Bayesian state space filter can be written as

$$\mathcal{C}(p(\mathbf{x}_{n+1}, \mathcal{A}_{n+1} | \tilde{\mathbf{r}}_{n+1}(t, \mathbf{c}_n))) = \bar{g}(\epsilon_{n+1|n+1}(\mathbf{c}_n)), \quad (19)$$

where $\epsilon_{n+1|n+1}(\mathbf{c}_n)$ is the predicted posterior state-estimation error depend on the control parameters and $\bar{g}(\cdot)$ defines the cost-to-go function of the transmitter. The conditional entropy was discussed as a possible information measure for the feedback, thus a possible cost-to-go function $\bar{g}(\cdot)$ of the transmitter is the conditional entropy of the predicted posterior state-estimation error $\epsilon_{n+1|n+1}(\mathbf{c}_n)$, given as $\bar{g}(\epsilon_{n+1|n+1}(\mathbf{c}_n)) = h(\epsilon_{n+1|n+1}(\mathbf{c}_n))$ [22], [49]. This entropy conditioned on the control parameter vector \mathbf{c}_n is directly coupled with the posterior covariance matrix of the Bayesian tracking filter that is lower bounded by the inverse of the EFIM in (4). The entropy of the predicted posterior state-estimation error (when assuming a Gaussian approximation) is given as

$$h(\epsilon_{n+1|n+1}(\mathbf{c}_n)) \propto \det(\tilde{\mathbf{C}}_{\mathbf{x}_{n+1}}(\mathbf{c}_n)), \quad (20)$$

where $\tilde{\mathbf{C}}_{\mathbf{x}_{n+1}}(\mathbf{c}_n)$ and $\mathbf{I}_{\mathbf{x}_{n+1}}(\mathbf{c}_n)$ is the predicted state covariance matrix as described in Section V-A of the state vector

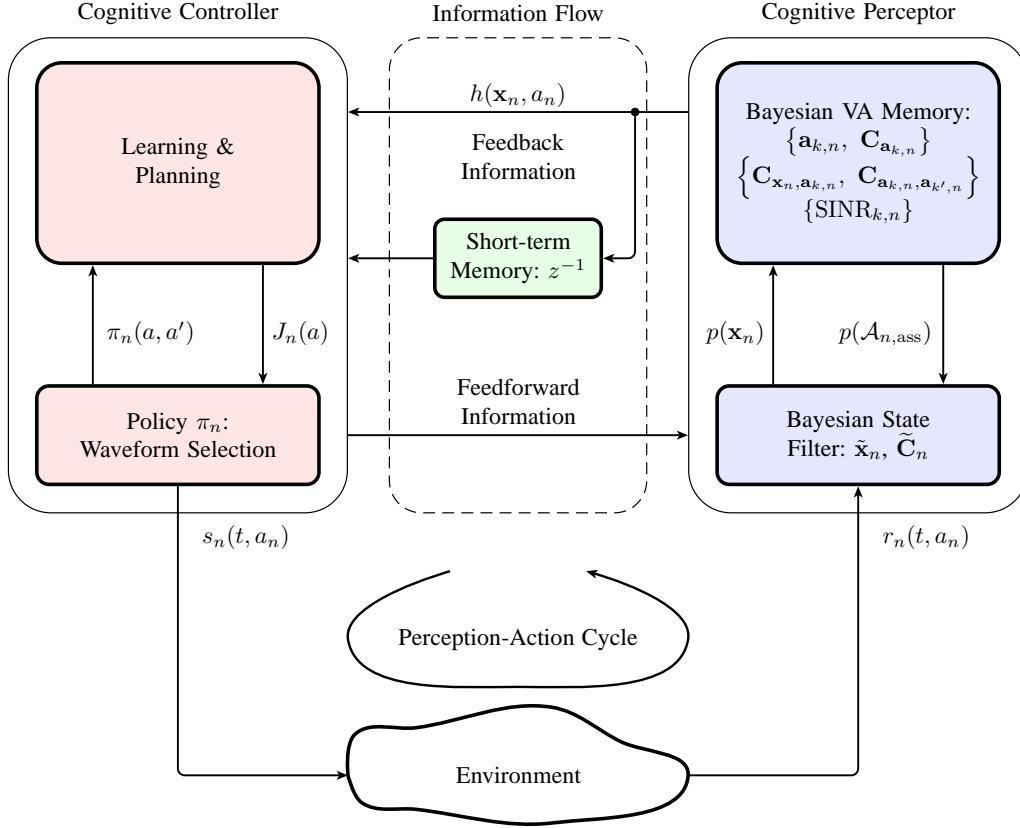


Fig. 3. Block diagram of the cognitive indoor positioning and tracking system that uses multipath channel information.

provided from the Bayesian state space filter (UKF) dependent on the control parameter vector \mathbf{c}_n . Thus, the entropy in (20) is directly coupled with the position-related information that is contained in the measurement noise covariance matrix $\mathbf{R}_{z,n}$ described by (18). How the introduced algorithm is using the state space and measurement model equations of the Bayesian state space estimator is described in more detail in Sections VII-B and VIII.

VI. REINFORCEMENT LEARNING (RL)

For readability of the following derivations of the control optimization algorithm, the cost-to-go of the CC (20) is rewritten as $\bar{g}(\epsilon_{n+1|n+1}(\mathbf{c}_n)) = h(\mathbf{x}_{n+1}, \mathbf{c}_n)$ with $\mathbf{c}_n \in \mathcal{A}$, where \mathcal{A} is the space of cognitive action with size $|\mathcal{A}|$ that represents the waveform library in our case. Consequently, the next set of waveform parameters has to be chosen in order to minimize the cost-to-go of the next posterior entropy. As elaborated in [50], dynamic programming represents an optimal solution for such problems, but unfortunately it is based on the assumption that the state to be controlled is “perfectly” perceivable. Hence, methods have been introduced that are capable of handling imperfect state information [51] with the drawback that they are computationally complex. In [5], [49] approximate dynamic programming was used for optimal control. In there, the trace of the posterior covariance matrix

was used as cost-to-go function to reduce the computational complexity. The policy for control parameter selection in the transmitter at time instance n is seeking to find the set of waveform parameters, for which the cost-to-go function $\bar{g}(\epsilon_{n+1|n+1}(\mathbf{c}_n)) \approx \text{tr}[\tilde{\mathbf{C}}_{\mathbf{x}_{n+1}}(\mathbf{c}_n)]$ is minimized for a rolling future horizon of l_{future} predicted states. In practice, it is difficult to construct all state transition probabilities from one state to another that are conditioned on the selected actions, including their cost incurred as a result of each transition. RL³ [52] represents an approximation of dynamic programming [50], [51] for solving such optimal control and future planning task with high computationally efficiency. In RL literature the cost-to-go function is termed value-to-go function $J_n(\mathbf{c}_n)$ that is updated online for every PAC based on the immediate rewards r_n . The immediate reward r_n is a measure of “quality” of an action \mathbf{c}_n taken on the environment. Using the Markovian assumption and following the way in [7], it is given by

$$r_n = g_n(h(\mathbf{x}_{n-1}, \mathbf{c}_{n-1}) - h(\mathbf{x}_n, \mathbf{c}_n)), \quad (21)$$

where $h(\mathbf{x}_n, \mathbf{c}_n) \propto \det(\mathbf{C}_{\mathbf{x}_n}(\mathbf{c}_n))$ and $g_n(\cdot)$ is an arbitrary scalar operator that in its most general form could also depend on the time instance n [7]. A reasonable function for the

³RL represents an intermediate learning procedure that lies between supervised and unsupervised learning as stated in [49].

reward is the scaled change in the posterior entropy from one PAC to the next, i.e.

$$r_n = \text{sign}(\Delta h(\mathbf{x}_n, \mathbf{c}_n)) |\log(|\Delta h(\mathbf{x}_n, \mathbf{c}_n)|)|. \quad (22)$$

A positive reward will be favoring the current action a_n for the future action \mathbf{c}_{n+1} and conversely a negative one will lead to a penalty for these actions. As described in [7], the cognitive RL algorithm has to find the optimal future action \mathbf{c}_{n+1} for the next PAC based on the immediate reward r_n and the learned value-to-go function $J_n(\mathbf{c}_n)$.

For computing the expected costs of future actions as it is done in dynamic programming, RL divides the computation of the value-to-go function into two parts, (i) the learning phase that incorporates the actual measured reward into the value-to-go function based on actions \mathbf{c}_n and \mathbf{c}_{n-1} , and (ii) the planning phase that incorporates predicted future rewards into the value-to-go function. Whereas for learning a “real” reward is perceived from the environment, for planning just model-based predicted rewards are perceived from the internal perceptor memory using the feedforward link. A faster convergence to the optimal control policy can be achieved in this way.

VII. LEARNING AND PLANNING: ALGORITHM

The value-to-go function that is used in the cognitive controller is defined as [7]

$$J_n(\mathbf{c}) = \mathbb{E}_{\pi_n} \{r_n + \gamma r_{n+1} + \gamma^2 r_{n+2} + \dots | \mathbf{c}_n = \mathbf{c}\}, \quad (23)$$

where r_n with $\mathbf{c} \in \mathcal{C}$ is the actual reward, r_{n+l} are the predicted future rewards that are based on the GPEM and GSCM parameter that are used by the Bayesian filter, $0 < \gamma \leq 1$ is the discount factor for future rewards based on action $\mathbf{c}_n \in \mathcal{C}$ and the expected value is calculated using the cognitive policy

$$\pi_n(\mathbf{c}', \mathbf{c}) = \mathbb{P}[\mathbf{c}_{n+1} = \mathbf{c}' | \mathbf{c}_n = \mathbf{c}], \quad \mathbf{c}, \mathbf{c}' \in \mathcal{C}, \quad (24)$$

where $\mathbb{P}[\cdot|\cdot]$ defines a conditional PMF that describes the transition probabilities of all actions $\mathbf{c} \in \mathcal{C}$ over time instances n . Following the derivations in [7], the value-to-go function can be reformulated in an incremental recursive manner, yielding

$$J_n(\mathbf{c}) \leftarrow J_n(\mathbf{c}) + \alpha \left[\mathcal{R}(\mathbf{c}) + \gamma \sum_{\mathbf{c}'} \pi_n(\mathbf{c}, \mathbf{c}') J_n(\mathbf{c}') - J_n(\mathbf{c}) \right], \quad (25)$$

where $\mathcal{R}(\mathbf{c}) = \mathbb{E}_{\pi_n} \{r_n | \mathbf{c}_n = \mathbf{c}\} \forall \mathbf{c} \in \mathcal{C}$ denotes the expected immediate reward and $\alpha > 0$ is the learning rate. The algorithm for updating the value-to-go function can be found in the Appendix of [4]. The incremental recursive update in (25) means that for all actions $\mathbf{c} \in \mathcal{C}$ the value-to-go function is updated using the expected immediate reward and the policy $\pi_n(\mathbf{c}, \mathbf{c}')$ for all these actions.

A. Learning from applied Actions

With the value of the immediate reward r_n , a new value is learned for the value-to-go function for the currently selected action \mathbf{c}_n using $J_n(\mathbf{c}_n) \leftarrow (1 - \alpha)J_n(\mathbf{c}_n) + \alpha \mathcal{R}(\mathbf{c}_n)$ of (25). This accounts for the “real” physical action on the environment. Hence, only one parameter set can be chosen

as an action for the PAC at a time; it would take at least $|\mathcal{C}|T$ seconds for applying all actions on the environment and collecting the according immediate rewards, where T is the time period of a PAC. Unfortunately, this results in a poor convergence rate of the algorithm and unacceptable behavior for time-variant environments. A possible remedy against this is the planning of future actions based on the state space and measurement model of the Bayesian state estimator.

B. Planning for Improving Convergence Behavior

Planning is defined as predicting expected future rewards using the state and measurement model of the Bayesian state space filter to improve the convergence rate of the RL algorithm. As depicted in Fig. 3, the feedforward link is used to connect the controller with the perceptor. The feedforward information is a hypothesized future action, which is selected for a future planning stage. Inspecting (25), one can observe that for every action $\mathbf{c} \in \mathcal{M}$, where $\mathcal{M} \subset \mathcal{C}$ is a subset of \mathcal{C} depending on the actual policy π_n , the predicted posterior covariance matrices $\hat{\mathbf{C}}_{n+l}(\mathbf{c})$ and the according predicted future rewards r_{n+l} , are computed with decreasing discount factor γ^l for predicted future rewards, for $l = 1, \dots, l_{\text{future}}$, where l_{future} is the future horizon. The predicted covariance matrices $\hat{\mathbf{C}}_{n+l}(\mathbf{c})$ for a specific future action \mathbf{c} is computed using the state space (e.g. (16)) and measurement model (e.g. (17)) of the Bayesian state space estimator and the according GPEM and GSCM parameters stored in the perceptors’ memory as shown in Fig. 3. After the planning process is finished, the value-to-go function is updated for all actions $\mathbf{c} \in \mathcal{M}$. Finally, the actual PAC is closed by updating the policy to π_{n+1} using the value-to-go function J_{n+1} and choosing the new action, i.e. the waveform parameters, for the next PAC according to this new policy. This means that the value-to-go function $J_n(\mathbf{c}_n)$ and the policy π_n are updated iteratively from one another from one PAC to the next PAC, with one important detail which is discussed below.

a) *Explore/Exploit trade-off*:: Both the planning process and choosing new actions are based on the policy. In planning, the chosen action-subset \mathcal{M} is defined by sampling from the policy π_n and new actions are selected based on the updated policy π_{n+1} . Hence, the policy is responsible for the explore/exploit trade-off in the action space. A widely used method for balancing the exploration of new actions and exploiting the already learned value-to-go function $J_n(\mathbf{c}_n)$ is the ϵ -greedy strategy, meaning that with a small probability of ϵ a random action is selected, representing pure exploration, and with probability of $1 - \epsilon$ the action is chosen according to the maximum of the value-to-go function, representing a pure exploitation. The random selection of a new action and the action in the subspace \mathcal{M} can either be selected from a uniform distribution over the action space \mathcal{A} or from the policy π_n . The policy is computed using the Boltzmann distribution

$$\pi_{n+l} = \pi_{n+l-1}(\mathbf{c}) \frac{\exp\{\Delta J_{n+l}(\mathbf{c})/\tau\}}{\sum_{\mathbf{c}'} \pi_{n+l-1}(\mathbf{c}') \exp\{\Delta J_{n+l}(\mathbf{c}')/\tau\}},$$

where τ defines the exploration degree and is referred to as the system temperature [53] and $\Delta J_{n+l}(\mathbf{c}) = J_{n+l}(\mathbf{c}) -$

$J_{n-1+l}(\mathbf{c})$. The cognitive action is selected according to

$$\mathbf{c}_n = \begin{cases} \text{random action } \sim \pi_{n+1} \in \mathcal{C} & \text{if } \xi < \epsilon \\ \arg \max_{\mathbf{c} \in \mathcal{C}} J_n(\mathbf{c}) & \text{otherwise} \end{cases}, \quad (26)$$

where $0 \leq \xi \leq 1$ is a uniform random number drawn at each time step n . As we have said, from the policy in (26) the new action \mathbf{c}_{n+1} is selected and applied on the environment so that the next PAC can start. The important concept of *attention* at the perceptor as well as the actuator side in the cognitive dynamic system can be argued with the following:

- **Perceptual attention:** Is given by the fact that the environment dependent parameters, i.e. the marginal PDF of the VA $p(\mathbf{a}_{k,n})$ and their multipath channel dependent reliability measures, $\text{SINR}_{k,n}$, are learned and updated online, so that the perceptual Bayesian state space filter puts its attention on the relevant position-related information in the received signal.
- **Control attention:** Is given by the fact that the policy π_n that is learned over time and the according subset of actions \mathcal{M} put focus on the “more relevant” actions. These action in turn focus on the relevant position-related information in the received signal.

VIII. WAVEFORM LIBRARY

The general form of the waveform library contains the control parameters $\mathbf{c}_n = \{T_{p,n}^{(j)}, f_{c,n}^{(j)}\}_{j=1}^J$ for the j -th anchor consisting of carrier frequencies and pulse durations. Hence, the VA specific MPC parameters are estimated using specific sub-bands of the radio channel spectrum defined by the parameter pair $T_{p,n}^{(j)}$ and $f_{c,n}^{(j)}$, which in turn is chosen in an “optimal” manner. Optimal in this case means that the position-related information that is contained in the MPC parameters is maximized at agent position \mathbf{p}_n (see for (4)).

Equations (5) and (4), which describe the parameters $\widetilde{\text{SINR}}_{k,n}^{(j)}$, show the relation between the pulse parameter pair $T_{p,n}^{(j)}$ and $f_{c,n}^{(j)}$ and the position-related information contained in the channel. The pulse duration $T_{p,n}^{(j)}$ scales the amount of DM and is directly proportional to the effective root mean square bandwidth β . The relation to $f_{c,n}^{(j)}$ is not that obvious, since it describes the frequency dependency of the environment parameters and thus the GSCM parameters as the complex amplitudes of the MPC and the DM PDP. The set of selected VA should lead to the highest overall SINR values (and accordingly the smallest range variances $\text{var}\{\hat{d}_{k,n}^{(j)}\}$) and the smallest possible GDOP⁴, i.e. geometric optimal constellation of VA positions which is reflected by the ranging direction matrix $\mathcal{I}_r(\phi_{k,n}^{(j)})$. In a cognitive sense this means that the actions $a \in \mathcal{A}$ are chosen to reduce the posterior entropy over time under quasi-stationary environment conditions.




Fig. 4. Scenario for probabilistic MINT using cognitive sensing in presence of additional DM interference. The anchors are at the positions $\mathbf{a}_1^{(1)}$ and $\mathbf{a}_1^{(2)}$. The black line represents the agent trajectory and the red part of the line indicates the agent positions, where the DM interference is activated.

IX. RESULTS

A. Measurement Setup

For the evaluation of this positioning approach, we use the seminar room scenario of the MeasureMINT database [55]. The measurements allow for 5 trajectories consisting of 1000 agent positions with a 1 cm spacing as shown in Fig. 4. At each position, UWB measurements are available of the channel between the agent and the two anchors at the positions $\mathbf{a}_1^{(1)} = [0.5, 7]^T$ and $\mathbf{a}_1^{(2)} = [5.2, 3.2]^T$. The measurements have been performed using an M-sequence correlative channel sounder developed by *Ilmsens*. This sounder provides measurements over approximately the FCC frequency range, from 3 – 10 GHz. On anchor and agent sides, dipole-like antennas made of Euro-cent coins have been used. They have an approximately uniform radiation pattern in azimuth plane and zeroes in the directions of floor and ceiling.

The chosen initial pulse duration is $T_p = 0.5$ ns (corresponding to a bandwidth of 2 GHz) and the center frequency is $f_c = 7$ GHz. The VA for the anchors at the positions $\mathbf{a}_1^{(1)}$ and $\mathbf{a}_1^{(2)}$ were computed a-priori up to order 2. The past window of agent positions for the SINR estimation is again chosen to be $w_{\text{past}} = 40$. For all simulations 30 Monte Carlo runs were conducted.

⁴The GDOP the ratio between position variance and the range variance [54]. For positioning a small value indicates a high level of confidence that high precision can be reached. Hence, the GDOP indicates a “good” geometry for positioning, i.e. a good geometric placement of the anchors.




Fig. 5. Performance CDF of the cognitive MINT algorithm using a smaller restricted set of VA. Visibilities of VA are computed using the SINR instead of optical ray-tracing.

B. Initial Experiment Setup

For the sake of simplicity, we reduce the control parameters to just the carrier frequency $c_n = f_{c,n}$ for each PAC for all anchors and we fix the pulse duration T_p . This means that the cognitive MINT system adaptively finds the carrier frequency $f_{c,n}$ from PAC to PAC that yields the highest reward from the environment by maximizing the position-related information. Starting from the initial value $f_{c,1} = 7$ GHz (which represents the center of the measured bandwidth), the carrier frequency is adapted over time using the posterior entropy in (20).

The finite space of cognitive actions \mathcal{C} contains the discrete frequency values bounded by the measured bandwidth, i.e. $f_{c,n,i} \in \mathcal{C}$, where $i = 1, \dots, |\mathcal{C}|$. The frequency spacing between the frequency bins is equidistant, $\Delta f_c = f_{c,n,i+1} - f_{c,n,i}$. For the experiments, we have chosen $\Delta f_c = 50$ MHz, considering the large signal bandwidth of 2 GHz. The starting policy is defined as a uniform distribution $\pi_1(\mathbf{c}', \mathbf{c}) = \mathcal{U}(f_{c,n,1}, f_{c,n,|\mathcal{C}|})$ and the cost-to-go function is chosen to be $J_1(\mathbf{c}) = 0 \forall \mathbf{c}$. The size of the planning subspace is $|\mathcal{M}| = 20$; the size of \mathcal{C} is $|\mathcal{C}| = 40$.

C. Discussion of Performance Results

1) *Conventional MINT*: Fig. 5 shows the overall position error CDF for “conventional” MINT (which assumes perfect floor plan knowledge) with and without cognitive waveform adaptation. To show the advantage of the cognitive MINT algorithm, a restricted set of VA is chosen and the visibilities of the VA are computed using the SINR instead of optical ray-tracing. As the CDF of “conventional” MINT indicates (blue line with circle marker), the tracking algorithm tends to diverge since too little position-related information is available. The black and the red lines show the overall position error CDF for cognitive MINT for a future horizon window of $l = 1$ and $l = 5$, respectively. As one can observe, the performance is significantly increased due to the cognitive waveform adaptation. This means that the cognitive MINT algorithm is able to increase the amount of position-related information by changing the sensing spectrum via the carrier frequency $f_{c,n,i} \in \mathcal{A}$ to bands that carry more geometry-dependent information in the MPC. Another interesting observation of Fig. 5 is that an increase of the planning horizon results in an




Fig. 6. Performance CDF of cognitive probabilistic MINT using a smaller restricted set of VA. For probabilistic MINT, the visibilities of VA are always computed using the SINR.


increased performance, confirming the correct functionality of the cognitive algorithm.

2) *Probabilistic MINT*: Fig. 6 shows the overall position error CDF for probabilistic MINT with and without cognitive waveform adaptation. Uncertainties in the floor plan and wrong associations can be robustly handled due to the probabilistic treatment of VA and thus none of the individual trajectory runs diverges. The already achieved high accuracy and robustness of probabilistic MINT are the reasons that cognitive sensing leads to only a minor additional performance gain for this scenario. It is suspected that for lower bandwidth the performance gain induced by the cognitive probabilistic MINT should be much more distinct.


3) *Probabilistic MINT with additional DM Interference*: In the last setup, we additionally have added synthetic DM interference filtered at a carrier frequency $f_c = 7$ GHz, with a bandwidth of 2 GHz. The DM parameters are chosen according to [56] except for the DM power. The experiments were conducted with three levels of DM power, $\Omega_1 = 1.1615 \cdot 10^{-9}$, $\Omega_1 = 5.8076 \cdot 10^{-9}$ and $\Omega_1 = 1.1615 \cdot 10^{-8}$.

Fig. 4 illustrates the scenario used for the experiment. The black line represents the agent trajectory and the red part of it indicates the agent positions, where the DM interference is activated. Fig. 7 shows the signals exchanged between the agent and the Anchors 1 and 2 for one sample position. The “clean” signals are shown in Fig. 7a, the noisy signal for DM power of $\Omega_1 = 1.1615 \cdot 10^{-9}$ in Fig. 7b. Looking at Fig. 7b it is quite obvious that this level of DM represents already a severe interference. The justification of using such a interference noise model lies in the fact that it can describe many kinds of measurement modeling mismatches, e.g. if the anisotropy of the antenna pattern for different angle of arrivals is not considered.

Fig. 8 illustrates the mean values of the cognitively adapted carrier frequency along one of the trajectories at DM power $\Omega_1 = 1.1615 \cdot 10^{-8}$. The mean is computed using the 30 Monte Carlo simulations of the experiment. The black line denotes the initial carrier frequency $f_{c,1}$ and the blue one the mean of the cognitively adapted carrier $f_{c,n}$. The blue dashed lines show a few example realizations of cognitively adapted carrier frequencies along different trajectories and for different



(a) Clean signals



(b) Noisy signal with DM power of $\Omega_1 = 1.1615 * 10^{-8}$

Fig. 7. Signals exchanged between agent and Anchors 1 and 2 for an example agent position. The gray lines represent the estimated delays of the MPC. Fig. 7a shows the “clean” signal and Fig. 7b the noisy signal.




Fig. 8. Mean carrier frequency for DM power $\Omega_1 = 1.1615 * 10^{-8}$. The black line denotes for the initial carrier frequency $f_{c,0}$ and the blue one the mean of the cognitively adapted carrier frequency $f_{c,n}$. The blue dashed lines show a few example realizations of cognitively adapted carrier frequencies along different trajectories and for different Monte Carlo runs.

Monte Carlo runs. The figure shows quite clearly that the cognitive probabilistic MINT algorithm is avoiding (almost at all agent positions, where additional DM interference is present) carrier frequencies $f_{c,n}$ near to the carrier of DM.

Fig. 9 shows the according mean entropy values of proba-




Fig. 9. Mean entropy of probabilistic MINT and cognitive probabilistic MINT over time instances n for DM power $\Omega_1 = 1.1615 * 10^{-8}$. The red and black dashed lines show a few example entropy realizations along different trajectories and for different Monte Carlo runs.

bilistic MINT (red line with diamond markers) and cognitive probabilistic MINT (black line with triangle markers) over time instances n for DM power $\Omega_1 = 1.1615 * 10^{-8}$. The red and black dashed lines show a few example entropy realizations along different trajectories and for different Monte Carlo runs. Before the noise disturbance starts the entropy of the probabilistic MINT algorithm is almost the same as of the cognitive probabilistic MINT algorithm. In the moment the disturbance is introduced, the entropy of the posterior increases. The cognitive probabilistic MINT algorithm then starts to change its carrier frequency $f_{c,n}$ (as shown in Fig. 8) until the entropy is again reduced. This leads to an almost constant or even decreasing entropy even in the presence of a tremendous noise level (black line with triangle markers in Fig. 9). In contrast to that the probabilistic MINT algorithm without cognitive waveform adaptation starts to diverge after the disturbance is introduced and is not able to recover. This is indicated by the rapid increase of the entropy and stagnation at a large value shown in Fig. 9 by the red line with diamond markers.

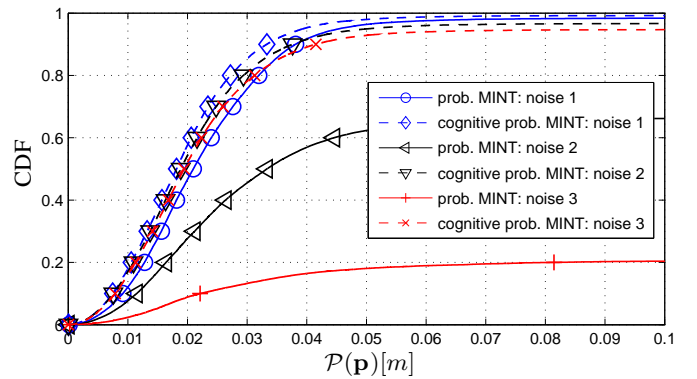


Fig. 10. Performance CDF of the cognitive probabilistic MINT algorithm with introducing a disturbance at three different noise levels along a certain part of the trajectory. Noise 1 corresponds to DM with $\Omega_1 = 1.1615 * 10^{-9}$, Noise 2 with power $\Omega_1 = 5.8076 * 10^{-9}$ and with power $\Omega_1 = 1.1615 * 10^{-8}$

This result is confirmed by looking at the performance CDF of the agent position error shown in Fig. 10. This comparison between probabilistic MINT and cognitive probabilistic MINT

illustrates the powerful property of the cognitive algorithm to separate relevant from irrelevant information using adaptation of the control parameter $f_{c,n}$ to avoid the noisy frequency band of the signal. The probabilistic MINT algorithm without waveform adaptation tends to diverge under such harsh conditions as depicted by CDF drawn with solid lines. In contrast to this, the cognitive MINT algorithm overcomes these impairments, leading again to a *robust* behavior as depicted by CDF drawn with dashed lines.

REFERENCES

- [1] J. M. Fuster, *Cortex and Mind - Unifying Cognition*. Oxford University Press, 2003.
- [2] P. Meissner, "Multipath-Assisted Indoor Positioning," Ph.D. dissertation, Graz University of Technology, 2014.
- [3] E. Leitinger, P. Meissner, C. Rudisser, G. Dumphart, and K. Witrisal, "Evaluation of Position-Related Information in Multipath Components for Indoor Positioning," *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 11, pp. 2313–2328, Nov 2015.
- [4] E. Leitinger, "Cognitive Indoor Positioning and Tracking using Multipath Channel Information," Ph.D. dissertation, Graz University of Technology, 2016.
- [5] S. Haykin, Y. Xue, and P. Setoodeh, "Cognitive Radar: Step Toward Bridging the Gap Between Neuroscience and Engineering," *Proceedings of the IEEE*, vol. 100, no. 11, pp. 3102–3130, nov. 2012.
- [6] S. Haykin, M. Fatemi, P. Setoodeh, and Y. Xue, "Cognitive Control," *Proceedings of the IEEE*, vol. 100, no. 12, pp. 3156–3169, dec. 2012.
- [7] M. Fatemi and S. Haykin, "Cognitive Control: Theory and Application," *Access, IEEE*, vol. 2, pp. 698–710, 2014.
- [8] A. Amiri and S. Haykin, "Improved Sparse Coding Under the Influence of Perceptual Attention," *Neural Comput.*, vol. 26, no. 2, pp. 377–420, Feb. 2014.
- [9] S. Haykin and J. Fuster, "On Cognitive Dynamic Systems: Cognitive Neuroscience and Engineering Learning From Each Other," *Proceedings of the IEEE*, vol. 102, no. 4, pp. 608–628, April 2014.
- [10] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.
- [11] P. Gregory, *Bayesian Logical Data Analysis for the Physical Sciences*. New York, NY, USA: Cambridge University Press, 2005.
- [12] D. S. Sivia and J. Skilling, *Data analysis : a Bayesian tutorial*, ser. Oxford science publications. Oxford, New York: Oxford University Press, 2006.
- [13] J. Borish, "Extension of the Image Model to arbitrary Polyhedra," *The Journal of the Acoustical Society of America*, March 1984.
- [14] J. Kunisch and J. Pamp, "An Ultra-Wideband space-variant Multipath Indoor Radio Channel Model," in *Ultra Wideband Systems and Technologies, 2003 IEEE Conference on*, Nov 2003, pp. 290–294.
- [15] E. Leitinger, P. Meissner, M. Lafer, and K. Witrisal, "Simultaneous Localization and Mapping using Multipath Channel Information," in *2015 IEEE International Conference on Communications Workshops (ICC)*, London, UK, June 2015, pp. 754–760.
- [16] D. Kershaw and R. Evans, "Optimal waveform selection for tracking systems," *Information Theory, IEEE Transactions on*, vol. 40, no. 5, pp. 1536–1550, sep 1994.
- [17] S. Haykin, A. Zia, Y. Xue, and I. Arasaratnam, "Control Theoretic Approach to Tracking Radar: First step towards cognition," *Digital Signal Processing*, vol. 21, no. 5, pp. 576–585, 2011.
- [18] K. Bell, C. Baker, G. Smith, J. Johnson, and M. Rangaswamy, "Cognitive radar framework for target detection and tracking," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 9, no. 8, pp. 1427–1439, Dec 2015.
- [19] G. Hoffmann and C. Tomlin, "Mobile Sensor Network Control Using Mutual Information Methods and Particle Filters," *Automatic Control, IEEE Transactions on*, vol. 55, no. 1, pp. 32–47, Jan 2010.
- [20] B. J. Julian, M. Angermann, M. Schwager, and D. Rus, "Distributed Robotic Sensor Networks: An Information-theoretic Approach," *Int. J. Rob. Res.*, vol. 31, no. 10, pp. 1134–1154, Sep. 2012. [Online]. Available: <http://dx.doi.org/10.1177/0278364912452675>
- [21] K. Chaloner and I. Verdinelli, "Bayesian Experimental Design: A Review," *Statistical Science*, vol. 10, no. 3, pp. 273–304, 1995. [Online]. Available: <http://www.jstor.org/stable/2246015>
- [22] T. M. Cover and J. A. Thomas, *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006.
- [23] H. L. Van Trees, *Detection, Estimation and Modulation, Part I*. Wiley Press, 1968.
- [24] S. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall Signal Processing Series, 1993.
- [25] I. I. Hussein, "Kalman filtering with optimal sensor motion planning," in *2008 American Control Conference*, June 2008, pp. 3548–3553.
- [26] M. Frohle, A. Zaidi, E. Strom, and H. Wymeersch, "Multi-step sensor selection with position uncertainty constraints," in *Globecom Workshops (GC Wkshps), 2014*, Dec 2014, pp. 1439–1444.
- [27] F. Bourgault, A. A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte, "Information based adaptive robotic exploration," in *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, vol. 1, 2002, pp. 540–545 vol.1.
- [28] T. Kollar and N. Roy, "Trajectory optimization using reinforcement learning for map exploration," *Int. J. Rob. Res.*, vol. 27, no. 2, pp. 175–196, Feb. 2008. [Online]. Available: <http://dx.doi.org/10.1177/0278364907087426>
- [29] —, "Efficient Optimization of Information-theoretic Exploration in SLAM," in *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 3*, ser. AAAI'08. AAAI Press, 2008, pp. 1369–1375. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1620270.1620287>
- [30] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT, 2006.
- [31] B. Grocholsky and B. Grocholsky, "Information-Theoretic Control of Multiple Sensor Platforms," Ph.D. dissertation, Department of Aerospace, Mechatronic and Mechanical Engineering, 2002.
- [32] J. L. William, "Information theoretic sensor management," Ph.D. dissertation, Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 2007.
- [33] M. P. Deisenroth, R. D. Turner, M. F. Huber, U. D. Hanebeck, and C. E. Rasmussen, "Robust filtering and smoothing with Gaussian processes," *IEEE Trans. Autom. Control*, vol. 57, no. 7, pp. 1865–1871, July 2012.
- [34] M. P. Deisenroth, D. Fox, and C. E. Rasmussen, "Gaussian processes for data-efficient learning in robotics and control," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 2, pp. 408–423, Feb 2015.
- [35] P. Meissner, E. Leitinger, and K. Witrisal, "UWB for robust indoor tracking: Weighting of multipath components for efficient estimation," *Wireless Communications Letters, IEEE*, vol. 3, no. 5, pp. 501–504, Oct 2014.
- [36] E. Leitinger, M. Fröhle, P. Meissner, and K. Witrisal, "Multipath-Assisted Maximum-Likelihood Indoor Positioning using UWB Signals," in *2014 IEEE International Conference on Communications Workshops (ICC)*, June 2014, pp. 170–175.
- [37] N. Michelusi, U. Mitra, A. Molisch, and M. Zorzi, "UWB Sparse/Diffuse Channels, Part I: Channel Models and Bayesian Estimators," *Signal Processing, IEEE Transactions on*, vol. 60, no. 10, pp. 5307–5319, 2012.
- [38] A. Molisch, "Ultra-Wide-Band Propagation Channels," *Proceedings of the IEEE*, vol. 97, no. 2, pp. 353–371, Feb. 2009.
- [39] Y. Shen and M. Win, "Fundamental Limits of Wideband Localization; Part I: A General Framework," *Information Theory, IEEE Transactions on*, vol. 56, no. 10, pp. 4956–4980, Oct. 2010.
- [40] Y. Shen, H. Wymeersch, and M. Win, "Fundamental Limits of Wideband Localization; Part II: Cooperative Networks," *Information Theory, IEEE Transactions on*, vol. 56, no. 10, pp. 4981–5000, Oct. 2010.
- [41] Y. Bar-Shalom and X.-R. Li, *Multitarget-Multisensor Tracking : Principles and Techniques*. Storrs, CT: Yaakov Bar-Shalom, 1995.
- [42] J. Vermaak, S. J. Godsill, and P. Perez, "Monte Carlo filtering for multi target tracking and data association," vol. 41, no. 1, pp. 309–332, Jan. 2005.
- [43] E. Leitinger, M. F., P. Meissner, K. Witrisal, and F. Hlawatsch, "Belief Propagation based Joint Probabilistic Data Association for Multipath-Assisted Indoor Navigation and Tracking," in *2016 International Conference on Localization and GNSS (ICL-GNSS)*, June 2016.
- [44] F. Meyer, P. Braca, P. Willett, and F. Hlawatsch, "Scalable Multitarget Tracking using Multiple Sensors: A belief propagation approach," in *Information Fusion (Fusion), 2015 18th International Conference on*, July 2015, pp. 1778–1785.
- [45] S. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*. Prentice Hall Signal Processing Series, 1998.
- [46] H. Wymeersch, "The Impact of Cooperative Localization on Achieving higher-level Goals," in *Communications Workshops (ICC), 2013 IEEE International Conference on*, June 2013, pp. 1–5.
- [47] F. Meyer, H. Wymeersch, M. Frohle, and F. Hlawatsch, "Distributed Estimation With Information-Seeking Control in Agent Networks," *Selected Areas in Communications, IEEE Journal on*, vol. 33, no. 11, pp. 2439–2456, Nov 2015.

- [48] C. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, The, vol. 27, no. 4, pp. 623–656, Oct 1948.
- [49] S. Haykin, *Cognitive Dynamic Systems: Perception-action Cycle, Radar and Radio*. New York, NY, USA: Cambridge University Press, 2012.
- [50] R. Bellman, *Dynamic Programming*. Princeton, NJ, USA: Princeton University Press, 1957.
- [51] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2000.
- [52] R. S. Sutton and A. G. Barto, *Introduction to Reinforcement Learning*, 1st ed. Cambridge, MA, USA: MIT Press, 1998.
- [53] A. Lazaric, M. Restelli, and A. Bonarini, "Reinforcement learning in continuous action spaces through sequential monte carlo methods," in *Advances in Neural Information Processing Systems*, 2007.
- [54] Z. Sahinoglu, S. Gezici, and I. Guvenc, *Ultra-wideband Positioning Systems – Theoretical Limits, Ranging Algorithms and Protocols*. Cambridge University Press, 2008.
- [55] P. Meissner, E. Leitinger, M. Lafer, and K. Witrals, "MeasureMINT UWB database," www.spsc.tugraz.at/tools/UWBmeasurements, 2013, Publicly available database of UWB indoor channel measurements. [Online]. Available: www.spsc.tugraz.at/tools/UWBmeasurements
- [56] J. Karedal, S. Wyne, P. Almers, F. Tufvesson, and A. Molisch, "A Measurement-Based Statistical Model for Industrial Ultra-Wideband Channels," *Wireless Communications, IEEE Transactions on*, vol. 6, no. 8, pp. 3028–3037, Aug. 2007.